

UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS FOR TRAFFIC SCHEDULING

Inventor:

Gregory G. Minshall

Prepared by:

Blakely, Sokoloff, Taylor & Zafman LLP
12400 Wilshire Blvd., Suite 700
Los Angeles, California 90025
(310)207-3800

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EL 485755535US Date of Deposit:

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service
"Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee
has been addressed to the Commissioner of Patents and Trademarks, Washington, D. C. 20231

Clara Wallin

(Typed or printed name of person mailing paper or fee)

Clara Wallin

(Signature of person mailing paper or fee)

6-20-01

(Date signed)

METHOD AND APPARATUS FOR TRAFFIC SCHEDULING

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The invention relates to the field of communications. More specifically, the invention relates to transmission over communication networks.

Background of the Invention

[0002] Various scheduling methods are used to support various levels of service. These services fall into one of two categories: priority based schedulers and round robin schedulers. A priority based scheduler always transmits the highest priority packets in one of its queues. A round robin scheduler transmits packets from each nonempty connection queue ("An Engineering Approach to Computer Networking", Keshav, p.236 (1997)). A weighted round robin scheduler transmits packets from each nonempty connection queue in proportion to each queue's "weight".

[0003] Unfortunately, higher priority packets can starve out lower priority traffic with a priority based scheduler. With a weighted round robin scheduler, low latency traffic may need to wait for an entire round (more than one packet transmission time) before being transmitted.

BRIEF DESCRIPTION OF THE DRAWING

[0004] Figure 1 is a diagram of a line card according to one embodiment of the invention.

[0005] Figure 2 is a diagram illustrating a packet scheduling mechanism according to one embodiment of the invention

[0006] Figure 3 is a flow chart for maintaining eligibility indicators according to one embodiment of the invention.

[0007] Figure 4 is a flow chart for a link scheduler according to one embodiment of the invention.

[0008] Figure 5 is a flowchart for a priority group scheduler according to one embodiment of the invention.

[0009] Figure 6 is a flowchart for a queue scheduler according to one embodiment of the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

[00010] In the following description, numerous specific details are set forth to provide a thorough understanding of the invention. However, it is understood that the invention may be practiced without these specific details. In other instances, well-known protocols, structures, processes and techniques have not been shown in detail in order not to obscure the invention.

[00011] Figure 1 is a diagram of a line card according to one embodiment of the invention. In Figure 1, the line card 101 is shown with a communication link 107. The communication link 107 connects to scheduler logic 109. The scheduler logic 109 connects to a set of queues 103. The set of queues 103 store traffic to be transmitted over the communication link 107. Although all of the traffic stored in the set of queues

103 will be transmitted over the communication link 107, the traffic can have different destinations. The queues 103 can correspond to customers, organizations, destinations, services, etc. The scheduler logic 109 determines when traffic stored in the queues 103 will be transmitted over the communications link 107.

[00012] Figure 2 is a diagram illustrating a packet scheduling mechanism according to one embodiment of the invention. In Figure 2, the set of queues 103 of Figure 1 store traffic. The set of queues 103 are configured into groups. Queues 211, 213, and 215 are configured as a first group 202. Queues 217, 219, A21, and A23 are configured as a second group 204. The queue A25 is configured as a third group 206. A queue scheduler 205 determines which queue in the first group of queues 202 will transmit traffic at a given time. A queue scheduler 207 determines which queue in the second group of queues 204 will transmit traffic. A queue scheduler 209 determines which queue in the third group of queues 206 will transmit traffic. A priority group scheduler 203 determines which of the group of queues 202, 204, or 206 will transmit traffic. A link scheduler 201 determines when a link associated with the set of queues 103 can transmit.

[00013] Figure 3 is a flow chart for maintaining eligibility indicators according to one embodiment of the invention. In Figure 3 at block 301, it is determined if a clock tick occurs. If a clock tick does not occur, then control loops back to block 301. If a clock tick occurs, then at block 302 a counter is incremented. At block 303, it is determined if the counter is equal to a link period. The link period can be adjusted in relation to the clock signals of a system. If the counter is not equal to or greater than the link period, then control flows back to block 301. If it is determined at block 303 that the counter is equal to or greater than the link period, then at block 305 a link balance is updated with a minimum of: 1) the link balance maximum; and 2) the link

balance incremented with the link token. At block 307, a priority group "clock" is updated with a priority group token. At block 309, the counter is reset. Control flows back to block 301 from block 309.

[00014] Figure 4 is a flow chart for a link scheduler according to one embodiment of the invention. At block 401 of Figure 4, it is determined if the hardware is ready to transmit. The hardware may be in a "flow controlled" or "not ready" state. The hardware may also still be transmitting the previous packet. If the hardware is not ready to transmit, then control loops back to block 401. If the hardware is ready to transmit, then at block 403 it is determined if the link is eligible to transmit. In one embodiment of the invention, if the link balance is equal to or less than a given value (e.g., zero), then the link is not eligible to transmit. If the link is not eligible to transmit, then control loops back to block 401, allowing another link to possibly transmit. In another embodiment of the invention, control does not loop back for another link because the links have individual flows occurring in parallel. If the link is eligible to transmit, then at block 405 it is determined if the link has data to transmit. If the link does not have data to transmit, then control loops back to block 401. In alternative embodiments of the invention, a "burst" value is maintained and updated when a link is eligible to transmit, but does not transmit. The burst value enables a link to transmit a burst of data after being idle. If the link has data to transmit, then at block 407, the link transmits data. At block 409, the link balance is updated. In one embodiment of the invention, the link balance is decremented by the cost of transmitting the data. The cost of the data transmission may bring the balance to zero or less than zero. In one embodiment of the invention, a lower limit is placed on the balance to prevent a link from being starved of transmission time after a large burst of data. The cost of the data can vary depending on implementation of the invention. The

cost of transmitting the data can be calculated based on the size of the data. The cost of transmitting the data could also be calculated using the size of the data and a modifier for the data type. These examples are intended to aid in understanding the invention and not meant to limit the invention.

[00015] Figure 5 is a flowchart for a priority group scheduler according to one embodiment of the invention. At block 501, a highest priority group is selected. At block 503, it is determined if the selected priority group has data to transmit. If the selected group does not have data to transmit, then at block 504 the scheduler determines if the selected priority group is the last priority group. If the selected priority group is not the last priority group, then the scheduler selects the next highest priority group at block 505. Control flows from block 505 to block 503. If the selected priority group has data to transmit, then at block 507 it is determined if the selected priority group is eligible to transmit. Eligibility can be determined in a number of ways as described above. In one embodiment of the invention, an eligibility value (initialized to zero) is compared with the priority group balance. If the eligibility value is less than the priority group balance, then the priority group is eligible to transmit data. If the selected priority group is eligible to transmit data, then at block 519 the data is transmitted from the selected priority group. At block D21, the eligibility value for the transmitting priority group is updated. From block D21, control flows to block 517 where the scheduler exits.

[00016] If the scheduler determines at block 507 that the selected priority group is not eligible, then at block 509 the scheduler determines if there is an ineligible higher priority group with data to transmit. If there is not an ineligible higher priority group with data, then at block 511 the selected priority group becomes a backup transmitting group. From block 511, control flows to block 504. If the scheduler determines at

block 509 that there is an ineligible higher priority group with data to transmit, then control flows to block 504. If the scheduler determines at block 504 that the selected priority group is the last priority group, then at block 513 the scheduler determines if there is a valid backup group. If the scheduler determines that there is not a valid backup group, then at block 517 the scheduler exits. If the scheduler determines at block 513 that there is a valid backup group, then at block 515 the backup group transmits its data. In another embodiment of the invention, ineligible priority groups are restricted from transmitting. Control flows from block 515 to block D21. In one embodiment of the invention, if an ineligible priority group transmits data, then the priority group balance is updated with the cost of the transmission.

[00017] Figure 6 is a flowchart for a queue scheduler according to one embodiment of the invention. At block 601, the queue scheduler determines which queues have data to transmit. At block 603, the queue scheduler selects the most eligible queue. The most eligible queue can be determined in a variety of ways. In one embodiment of the invention, the queue with the lowest eligibility value is the most eligible queue. In another embodiment of the invention, the queue with the highest eligibility value is the most eligible queue. In another embodiment of the invention, the queue with an eligibility value greater than all other eligibility values but less than a "clock" value is the most eligible queue. At block 605, data is transmitted from the selected queue. At block 607, an eligibility value for the selected queue is updated. In one embodiment of the invention, the eligibility value for the selected queue is used by the queue scheduler as a lower boundary for the next transmitting queue's eligibility value.

[00018] A packet scheduler combining features of priority-based schedulers and generalized processor sharing schedulers (i.e., round robin schedulers, fair queuing

schedulers, etc.) prevents higher priority traffic from starving lower priority traffic while preventing lower priority traffic from delaying higher priority traffic. Such a packet scheduler enables a network element to allocate different levels of service to different classes of traffic. A network element with this scheduler can offer various service levels including low-latency traffic, guaranteed bandwidth traffic, and best-effort traffic.

[00019] The described line card include memories, processors, and/or Application Specific Integrated Circuits ("ASICs"). Such memory includes a machine-readable medium on which is stored a set of instructions (i.e., software) embodying anyone, or all, of the methodologies described herein. Software can reside, completely or at least partially, within this memory and/or within the processor and/or ASICs. For the purpose of this specification, the term "machine-readable medium" shall be taken to include any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"), random access memory ("RAM"), magnetic disk storage media, optical storage media, flash memory devices, electrical, optical, acoustical, or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.), etc.

[00020] While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described.

[00021] The method and apparatus of the invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting on the invention.